

Learning Neural Operators from Partial Observations via Latent Autoregressive Modeling

Jingren Hou^{*1,3}, Hong Wang^{*2}, Pengyu Xu^{†1,3}, Chang Gao^{1,3}, Huafeng Liu^{1,3}, Liping Jing^{†1,3}

¹School of Computer Science and Technology, Beijing Jiaotong University, Beijing, China

²University of Science and Technology of China

³Beijing Key Laboratory of Traffic Data Mining and Embodied Intelligence, Beijing, China
{21112019, pengyuxu, 22110098, hfliu1, lpjing}@bjtu.edu.cn, wanghong1700@mail.ustc.edu.cn

Abstract

Real-world scientific applications frequently encounter incomplete observational data due to sensor limitations, geographic constraints, or measurement costs. Although neural operators significantly advanced PDE solving in terms of computational efficiency and accuracy, their underlying assumption of fully-observed spatial inputs severely restricts applicability in real-world applications. We introduce the first systematic framework for learning neural operators from partial observation. We identify and formalize two fundamental obstacles: (i) the supervision gap in unobserved regions that prevents effective learning of physical correlations, and (ii) the dynamic spatial mismatch between incomplete inputs and complete solution fields. Specifically, our proposed Latent Autoregressive Neural Operator (LANO) introduces two novel components designed explicitly to address the core difficulties of partial observations: (i) a mask-to-predict training strategy that creates artificial supervision by strategically masking observed regions, and (ii) a Physics-Aware Latent Propagator that reconstructs solutions through boundary-first autoregressive generation in latent space. Additionally, we develop POBench-PDE, a dedicated and comprehensive benchmark designed specifically for evaluating neural operators under partial observation conditions across three PDE-governed tasks. LANO achieves state-of-the-art performance with 18–69% relative L2 error reduction across all benchmarks under patch-wise missingness with less than 50% missing rate, including real-world climate prediction. Our approach effectively addresses practical scenarios involving up to 75% missing rate, to some extent bridging the existing gap between idealized research settings and the complexities of real-world scientific computing.

Code — <https://github.com/Kingyum-Hou/LANO>

1 Introduction

Neural operators provide a promising data-driven alternative for solving partial differential equations (PDEs) in science and engineering (Zachmanoglou and Thoe 1986). They learn to approximate the input-output mappings of PDE-governed

^{*}These authors contributed equally.

[†]Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

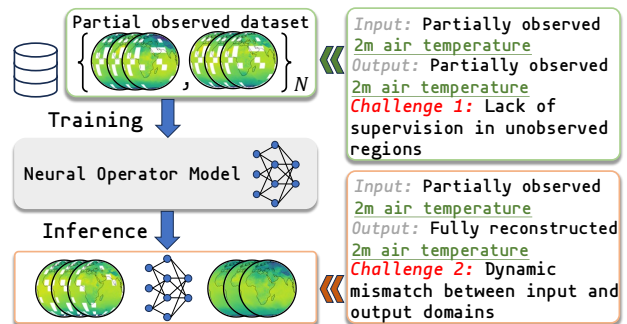


Figure 1: An overview of neural operator learning under partially observed PDE datasets. Once trained, the model is expected to infer solutions from unseen, partially observed inputs. However, two challenges arise in this setting, as summarized on the right.

tasks from data during training and then infer the PDE solutions for accurate simulations (Karniadakis et al. 2021; Li et al. 2020; Gao and Wang 2023). Compared to traditional PDE solvers, neural operators can infer solutions over five orders of magnitude faster during inference (Azizzadenesheli et al. 2024; Wang et al. 2024; Dong et al. 2024; Wang et al. 2025a). Beyond computational efficiency, neural operators successfully model a wide range of physical phenomena governed by underlying PDEs, including airflow, weather patterns, and optical systems (Li et al. 2023; Gupta and Brandstetter 2022; Brandstetter et al. 2022; Luo et al. 2024; Huang et al. 2025; Wang et al. 2025c; Lv, Liu, and Wang 2025; Wei et al. 2025b,a).

A major obstacle for neural operators is their dependency on high-quality and fully observed training datasets (Kovachki et al. 2023). However, in many real-world applications across the natural sciences and engineering, data are often partially observed—for instance, weather records miss large geographic areas where stations cannot be installed (Morice et al. 2021). Similar gaps arise in seismic exploration, electromagnetic surveys, and medical imaging (Zheng et al. 2019; Puzyrev 2019; Feng et al. 2021) due to terrain blockage, sparse sampling, or sensor failure. **Crucially, each training sample normally consists of an input field and its corre-**

sponding solution field on a fixed grid; Partial observation removes spatially correlated unobserved regions from both. These extensive and irregular unobserved regions make interpolation-based data completion unreliable for training neural operators.

This partial observability poses two challenges for operators learning (Fig. 1): (i) Lack of supervision in unobserved regions: In partial observations, the solution fields are only partially available. The absence of ground truth in unobserved regions hinders the model from learning intricate physical correlations of PDEs (Wu et al. 2024). (ii) Dynamic mismatch between partial input and global output domains: Unlike classical neural operator settings where inputs and outputs are defined over the same spatial grid (Brandstetter, Worrall, and Welling 2022; Wu et al. 2024), partially observed data results in mismatched and dynamically varying input regions. The model must infer full-domain outputs from partially observed inputs, which most existing architectures are not designed to handle (Gao et al. 2025).

Our approach To address these challenges, we propose a hybrid solution that combines a novel masking-based training strategy with a physics-aware autoregressive modeling framework, and we investigate: *can masked prediction tasks within observed regions implicitly guide operator networks to accurately infer unobserved regions?*.

Firstly, we introduce a mask-to-predict training strategy (MPT). Inspired by the masking strategy in natural language processing (NLP) and computer vision (CV), which is conceptually simple: it removes a portion of the data and learns to predict the removed content (Devlin et al. 2019; He et al. 2022), MPT artificially masks parts of the already observed input regions during supervised training. This creates pseudo-missing regions with available supervision, encouraging the model to extrapolate from partial context. Instead of pre-training on unlabeled data of neural operators (Chen et al. 2024), MPT is directly integrated into the training process, enabling better generalization to real unobserved regions.

Secondly, we propose the latent autoregressive neural operator (LANO). Directly predicting all unobserved regions simultaneously often results in blurry and incoherent solutions due to the inherent difficulty of modeling complex joint distributions over spatially dependent points. This issue has been similarly noted in visual content generation, where autoregressive or parallelized prediction strategies have effectively mitigated global modeling difficulties by explicitly handling local dependencies (Van Den Oord, Kalchbrenner, and Kavukcuoglu 2016; Wang et al. 2025d). Motivated by this observation, LANO employs an autoregressive generation strategy within a latent representation space, progressively reconstructing unobserved regions in a spatially structured manner. Central to this design is the Physics-aware Latent Propagator (PhLP), which decomposes solution prediction into a structured multi-step process. Starting from observed boundary conditions, PhLP incrementally propagates latent information into unobserved regions. This boundary-first design reflects the physical structure of many PDEs and enforces physical consistency in predictions.

To our knowledge, this is the first neural operator frame-

work that enables effective training from partially observed data with dynamic and unaligned input-output domains. Our main contributions are summarized as follows:

- We identify and formalize two key challenges in operator learning under partial observation: (i) the lack of supervision in unobserved regions, and (ii) the dynamic mismatch between partial input and global output domains
- We propose a hybrid framework to address these challenges: (i) the Mask-to-Predict (MPT) training strategy for robust learning from partial observation, and (ii) the Latent Autoregressive Neural Operator (LANO), which uses a physics-aware latent propagator (PhLP) to progressively reconstruct solutions across unobserved regions.
- We construct POBench-PDE, the first benchmark suite for operator learning under partial observation with six sets of comprehensive experiments across three PDE-governed tasks, and demonstrate that under patch-wise missingness with less than 50% missing rate, our model consistently outperforms existing methods by 17.8%-68.7% in relative error reduction on both synthetic and real-world PDE-governed tasks, including weather forecasting.

2 Related Work

Neural Operator Learning. This field aims to approximate PDE input-output mappings through deep models. DeepONet (Lu et al. 2021) pioneered this direction, followed by Fourier Neural Operator (FNO) (Li et al. 2020) and its variants (Li et al. 2024; Tran et al. 2023; Wen et al. 2022). To handle irregular meshes, Geo-FNO (Li et al. 2023) introduces coordinate mappings for non-uniform domains.

Transformer-based approaches have gained prominence, including GK-Transformer (Cao 2021), OFormer (Li, Meidani, and Farimani 2022), GNOT (Hao et al. 2023), ONO (Xiao et al. 2023), and MoE-POT (Wang et al. 2025b). Recent methods explore latent space modeling: Transolver (Wu et al. 2024) uses Physics-Attention for geometric features, LSM (Wu et al. 2023) operates in hierarchical latent space, while IPOT (Lee and Oh 2024), LNO (Wang and Wang 2024), and UPT (Alkin et al. 2024) decouple input-output sampling locations.

However, these methods assume shared input-output discretization or produce one-shot full-domain predictions, limiting their effectiveness under partial observation. In contrast, LANO enables progressive information propagation from observed to unobserved regions through boundary-first latent evolution.

Training under Sparse Observation. Some transformer-based methods (OFormer, IPOT, GNOT, LNO) handle partial inputs during inference but require fully observed training, limiting real-world application. DINO (Sitzmann et al. 2020) and CORAL (Serrano et al. 2023) employ implicit neural representations for continuous prediction from sparse signals.

These approaches either rely on implicit modeling or lack systematic training under dynamic sparse patterns, focusing primarily on sparse point-wise missingness rather than spatially correlated unobserved regions.

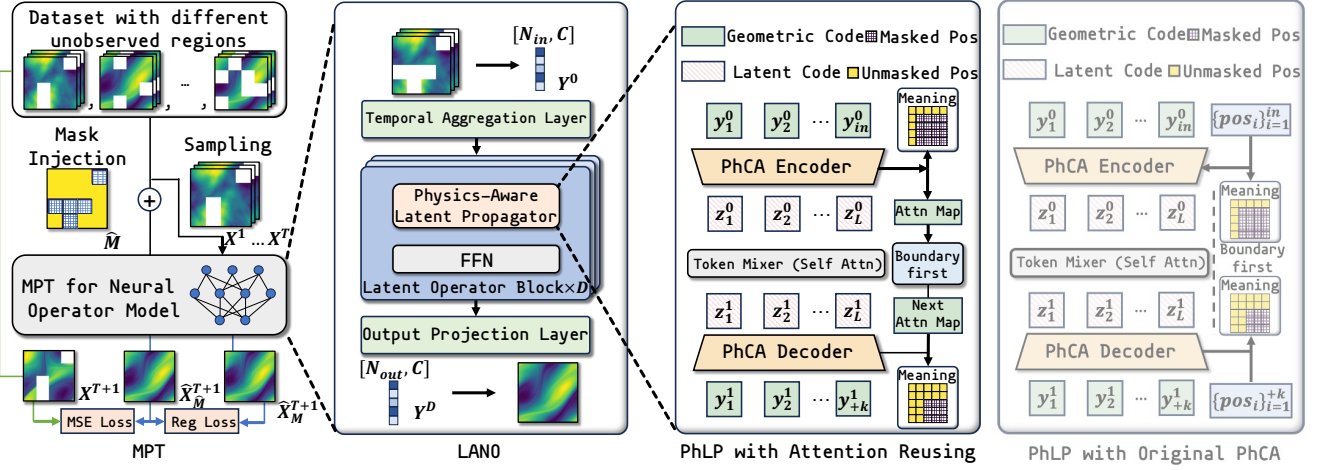


Figure 2: **Overview of the proposed LANO model architecture.** We begin by sampling trajectories from partial observations. The model is optimized using the mask-to-predict training strategy (MPT), where artificially masked past frames are used to predict the next frame. To support this learning objective, we design a novel model architecture incorporating a physics-aware latent propagator (PhLP), which implements a boundary-first autoregressive framework.

3 Proposed Method

3.1 Overview of LANO

Problem setup. We consider time-dependent PDEs defined on a spatial domain Ω , discretized into N points $\mathcal{D} \subset \mathbb{R}^{N \times d}$, where d is the spatial dimensionality. For a single trajectory, we denote the input and output (solution) measurements as $\mathbf{x}_\theta(\mathcal{D}_{in})$ and $\mathbf{x}_\theta(\mathcal{D}_{out})$, where $\mathcal{D}_{in}, \mathcal{D}_{out} \subset \mathcal{D}$ and the data lie in $\mathbb{R}^{N_{in} \times d}$ and $\mathbb{R}^{N_{out} \times d}$, respectively. In many real-world settings (e.g., weather forecasting), only trajectories of observations are available, while the underlying PDE parameters θ remain unknown. In this case, to predict the next timestep, the model must implicitly infer the most likely θ from adjacent T frames $\{\mathbf{x}^i\}_{i=1}^T$ (Hao et al. 2024). Our goal is to predict the solution at the next timestep given historical observations. For instance, in fluid dynamics, the input is a sequence of past velocity fields over \mathcal{D}_{in} , and the target is the full velocity field over \mathcal{D} . We also define a binary mask $M \in \{0, 1\}^N$ to indicate unobserved regions: $M_i = 1$ if position i is observed, and $M_i = 0$ otherwise.

Objective. We aim to learn a neural operator $\mathcal{G}_w(\mathbf{x}^{1:T}) = \mathbf{x}^{T+1}$ from partially observed temporal PDE datasets, where w denotes learnable parameters. The model autoregressively takes T frames as input and predicts the next frame. We directly supervise the one-step prediction loss using input-output pairs $\{\{\mathbf{x}^i\}_{i=1}^T(\mathcal{D}_o), \mathbf{x}^{T+1}(\mathcal{D}_o)\}_{\# \text{trajectory}}$, where $\mathcal{D}_o \subset \mathcal{D}$ denotes the observed regions with $N_o < N$ points, and each trajectory has a different \mathcal{D}_o ,

$$\min_w \mathcal{L} = \mathbb{E}_{\mathbf{x}} \sum_{T \leq i < T_{all}} \|\mathcal{G}_w(\mathbf{x}^{i-T+1:i} * M) - \mathbf{x}^{i+1}\|_2^2. \quad (1)$$

We propose a tailored training strategy to address the lack of supervision in unobserved regions. It injects artificial masks \hat{M} as structured noise during training, which helps mitigate cumulative error propagation in autoregressive neural PDE

solvers (Brandstetter, Worrall, and Welling 2022) and has been shown to effectively reduce long-term error accumulation in neural operators (Hao et al. 2024).

3.2 Mask-to-predict training strategy

To address the lack of supervision in unobserved regions, we propose a mask-to-predict training strategy (MPT) that integrates random input masking into neural operator training. Motivated by the observation that PDE dynamics are often invariant to sparsely sensed fields (Chen et al. 2024), MPT adapts masking-based learning from NLP and CV (Devlin et al. 2019; He et al. 2022) by injecting artificial masks into already incomplete inputs while retaining supervision over observed outputs.

Concretely, we apply a stochastic mask \hat{M} to further occlude input fields, training the model to recover solutions from limited context. This encourages extrapolation beyond available measurements and improves generalization to real-world unobserved regions. We additionally introduce consistency regularization that enforces prediction alignment between original and masked inputs, promoting invariance to artificial sparsity. More details are provided in the Appendix A.1.

3.3 Boundary-first autoregressive framework

Directly predicting all unobserved fields at once often results in blurry and inconsistent outputs, due to the complexity of spatial joint distributions. To mitigate this, we draw inspiration from parallel autoregressive paradigms in visual generation (Wang et al. 2025d), which model local dependencies progressively to improve stability and structure.

Given that PDE solutions are highly sensitive to boundary conditions, we adopt a boundary-first autoregressive framework that propagates information outward from known regions. To enhance efficiency, generation is performed in latent

space, which improves both computational cost and spatial consistency. This design allows the model to generate solutions in a physics-aware and coherent manner across partially observed domains.

3.4 Model Architecture

Overview. Existing neural operator architectures lack explicit positional guidance during feature extraction, limiting their predictive ability until the final projection stage. Inspired by LNO (Wang and Wang 2024), we propose a boundary-first autoregressive framework based on Physics-Cross-Attention (PhCA). This design introduces spatially aware information during feature propagation, enabling progressive prediction with enhanced physical and positional consistency.

Architecture. The proposed latent autoregressive neural operator (LANO) consists of three modules: a temporal aggregation layer for input embedding, latent operator layers that progressively propagate information from observed to unobserved regions guided by PhCA, and an output projection layer for geometric space recovery. The overall architecture is shown in Figure 2.

Temporal Aggregation Layer. This layer embeds input positions and physical measurements from T frames into deep features. Given positions set $\mathcal{D} = \{s_i\}_{i=1}^{N_{in}}$ with geometric information of N_{in} spatial points and observed fields $\{\mathbf{x}^i\}_{i=1}^T$, we embed them into features $\mathbf{Y}^0 = \{\mathbf{y}_i\}_{i=1}^{N_{in}}$ by a linear layer, where $\mathbf{y}_i \in \mathbb{R}^{1 \times C}$. The embedding is given by

$$\mathbf{Y}^0 = \text{Linear}(\text{Concat}(\mathcal{D}, \mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^T)). \quad (2)$$

Latent Operator Layer. This layer models PDE operators in the latent space to propagate features from input \mathbf{Y}^0 to output \mathbf{Y}^D using stacked latent operator blocks with a physics-aware latent propagator (PhLP),

$$\begin{aligned} \hat{\mathbf{Y}}^l &= \text{PhLP} \left(\text{LayerNorm} \left(\mathbf{Y}^{l-1} \right) \right) + \mathbf{Y}^{l-1} \\ \mathbf{Y}^l &= \text{MLP} \left(\text{LayerNorm} \left(\hat{\mathbf{Y}}^l \right) \right) + \hat{\mathbf{Y}}^l. \end{aligned} \quad (3)$$

Output Projection Layer adopts a linear projection upon the last deep feature \mathbf{Y}^D and obtains the output as a prediction of solution fields \mathbf{x}^{T+1} , the output can be calculated through $\mathbf{x}^{T+1} = \text{Linear}(\mathbf{Y}^D)$.

3.5 Physics-Aware Latent Propagator

To handle feature interactions across partially observed inputs, we introduce the Physics-Aware Latent Propagator (PhLP), which employs a boundary-first autoregressive framework to progressively reconstruct unobserved regions through physics-guided information propagation.

Physics-Cross-Attention Mechanism. PhLP operates through a Physics-Cross-Attention (PhCA) encoder-decoder in latent space. Given features $\mathbf{Y}^{l-1} \in \mathbb{R}^{N_{in} \times C}$ and observation mask $\mathbf{M} \in \{0, 1\}^{N_{in}}$, the encoder generates latent tokens $\mathbf{Z} \in \mathbb{R}^{H \times L \times C_h}$ through:

$$\begin{aligned} \mathbf{S} &= \text{softmax} \left(\frac{\text{MLP}(\mathbf{Y}_h)}{\tau} \right) \odot \mathbf{M} \\ \mathbf{Z} &= \frac{\mathbf{S}^T \mathbf{Y}_h}{\|\mathbf{S}\|_1 + \epsilon}, \end{aligned} \quad (4)$$

where $\mathbf{S} \in \mathbb{R}^{H \times N_{in} \times L}$ denotes attention maps aggregating spatial information into L latent tokens and $\mathbf{Y}_h \in \mathbb{R}^{H \times N_{in} \times C_h}$ is obtained by grouping the input features \mathbf{Y}^{l-1} along the channel dimension into H heads, each of dimension $C_h = C/H$. Notably, we directly utilize deep features as keys since they inherently encode spatial position information, eliminating the need for explicit positional encodings.

Boundary-First Propagation. PhLP employs partial convolution (PConv) (Liu et al. 2018) to propagate information from observed to unobserved regions:

$$\mathbf{S}^{next}, \mathbf{M}^{next} = \text{PConv}(\mathbf{S}, \mathbf{M}). \quad (5)$$

The updated tokens are processed through a self-attention before decoding:

$$\text{PhLP}(\mathbf{Y}^l) = \text{PhCA-Decoder}(\text{Attn}(\mathbf{Z}), \mathbf{S}^{next}). \quad (6)$$

For further details on PhLP, please refer to the Appendix A.2.

Theoretical Foundation. The solution of PDEs can be formulated as an iterative update process, with existing neural operator approaches leveraging Monte-Carlo sampling to approximate integral operators over the spatial domain \mathcal{D} for each update step (Li et al., 2020; Kovachki et al., 2023). To establish theoretical grounding for our method, we demonstrate that PhLP maintains equivalence to learnable integral operators defined on \mathcal{D} .

Theorem 3.1 (PhLP as an Integral Operator with Self-Update (Simplified)). *Let $\mathcal{D} \subset \mathbb{R}^d$ be the computational domain and $\mathcal{D}_o^l \subseteq \mathcal{D}$ the observed region at layer l . Given features $\mathbf{Y}^{l-1} : \mathcal{D} \rightarrow \mathbb{R}^C$, PhLP is equivalent to applying the learnable integral operator*

$$\mathcal{G}^l(\mathbf{Y}^{l-1})(\mathbf{x}^*) = \int_{\mathcal{D}_o^l} \kappa^l(\mathbf{x}^*, \boldsymbol{\xi}) \mathbf{Y}^{l-1}(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad \forall \mathbf{x}^* \in \mathcal{D} \quad (7)$$

where the kernel $\kappa^l(\mathbf{x}^*, \boldsymbol{\xi}) \approx \sum_{h=1}^H \sum_{k=1}^L \phi_{hk}^l(\mathbf{x}^*) \psi_{hk}^l(\boldsymbol{\xi})$ admits a low-rank factorization. For observed points $\mathbf{x}^* \in \mathcal{D}_o^l$, the kernel includes an identity contribution enabling residual self-update.

Proof Sketch. The proof proceeds through three key steps: (i) PhCA encoder produces attention weights $\mathbf{S} \in \mathbb{R}^{H \times N_o \times L}$ that define $\psi_{hk}^l(\boldsymbol{\xi}) := \mathbf{S}_{hk}(\boldsymbol{\xi})$; (ii) Decoder parameters provide $\phi_{hk}^l(\mathbf{x}^*)$, realizing the low-rank kernel through attention contraction $\mathbf{S}^T \mathbf{Y}$; (iii) The token mixer enriches the kernel through learnable token transformations while maintaining the operator form. See Appendix B for complete proof. \square

Variant: LANO-S. We investigate a variant (LANO-S) that uses explicit positional encodings to recalculate attention maps during decoding, following LNO’s approach and providing a direct comparison between implicit and explicit spatial representations.

MODEL	NAVIER-STOKES				DIFFUSION-REACTION				ERA5			
	POINT-WISE		PATCH-WISE		POINT-WISE		PATCH-WISE		POINT-WISE		PATCH-WISE	
TRAIN TEST	MISSING RATIO $s = 5\%$											
	5%	25%	5%	25%	5%	25%	5%	25%	5%	25%	5%	25%
MIONET (2022)	0.5662	/	0.6531	/	0.9244	/	0.9196	/	0.0487	/	0.0532	/
OFORMER (2022)	0.2020	0.2082	0.2021	0.2090	0.0301	0.0380	0.0334	0.0540	0.0332	0.0334	0.0289	0.0290
CORAL (2023)	0.2320	0.2510	0.2479	0.2515	0.4758	/	0.4916	/	/	/	/	/
GNOT (2023)	0.2311	0.2450	0.2574	0.2733	0.9166	0.9166	0.9235	0.9236	0.0256	0.0256	0.0255	0.0256
IPOT (2024)	0.2528	0.2526	0.2539	0.2556	<u>0.0230</u>	<u>0.0284</u>	<u>0.0319</u>	<u>0.0527</u>	0.0453	0.0458	0.0453	0.0453
LNO (2024)	0.1687	0.1745	0.1798	0.1879	/	/	/	/	0.0212	0.0213	0.0217	0.0216
LANO-S(OURS)	0.1649	0.1645	0.1621	0.1694	0.2081	0.4596	0.2344	0.4719	0.0168	0.0170	0.0157	0.0154
LANO (OURS)	0.1268	0.1275	0.1244	0.1310	0.0080	0.0089	0.0148	0.0281	0.0122	0.0123	0.0118	0.0120
Promotion	24.8%	26.9%	30.8%	30.3%	65.2%	68.7%	53.6%	46.7%	42.5%	42.3%	45.6%	44.4%
TRAIN TEST	MISSING RATIO $s = 25\%$											
	25%	50%	25%	50%	25%	50%	25%	50%	25%	50%	25%	50%
MIONET (2022)	0.5703	/	0.7312	/	0.9214	/	0.9262	/	0.0309	/	0.0460	/
OFORMER (2022)	0.2079	0.2159	0.2083	0.2226	0.0457	0.0670	0.0818	0.1569	0.0299	0.0302	0.0284	0.0288
CORAL (2023)	0.2264	0.2322	0.2480	0.2640	0.4796	/	0.5366	/	/	/	/	/
GNOT (2023)	0.2369	0.2527	0.2734	0.3009	0.9186	0.9186	0.9249	0.9255	0.0257	0.0257	0.0258	0.0258
IPOT (2024)	0.2568	0.2608	0.2556	0.2625	<u>0.0262</u>	<u>0.0364</u>	<u>0.0485</u>	<u>0.1181</u>	0.0452	0.0455	0.0453	0.0453
LNO (2024)	0.1732	0.1797	0.1915	0.2110	/	/	/	/	0.0213	0.0214	0.0214	0.0224
LANO-S(OURS)	0.1632	0.1687	0.1749	0.1935	0.4600	0.6504	0.4718	0.6931	0.0164	0.0164	0.0163	0.0167
LANO (OURS)	0.1274	0.1310	0.1435	0.1608	0.0092	0.0127	0.0275	0.0756	0.0121	0.0124	0.0120	0.0124
Promotion	26.4%	27.1%	25.1%	23.8%	64.9%	65.1%	43.3%	36.0%	43.2%	42.1%	43.9%	44.6%
TRAIN TEST	MISSING RATIO $s = 50\%$											
	50%	75%	50%	75%	50%	75%	50%	75%	50%	75%	50%	75%
MIONET (2022)	0.5754	/	0.8382	/	0.9651	/	0.9264	/	0.0402	/	0.0684	/
OFORMER (2022)	0.2151	0.2218	0.2348	0.2597	0.0972	0.1515	0.2843	0.4893	0.0283	0.0288	0.0273	0.0274
CORAL (2023)	0.2396	0.2662	0.2835	0.3414	0.4909	/	0.6225	/	/	/	/	/
GNOT (2023)	0.2807	0.3051	0.3210	0.3609	0.9138	0.9145	0.9267	0.9275	0.0256	0.0258	0.0320	0.0344
IPOT (2024)	0.2594	0.2645	0.2731	0.2878	<u>0.0405</u>	<u>0.0678</u>	<u>0.1212</u>	<u>0.3321</u>	0.0437	0.0438	0.0392	0.0389
LNO (2024)	0.1749	0.1921	0.2320	0.3434	/	/	/	/	0.0219	0.0220	0.0223	0.0276
LANO-S(OURS)	0.1601	0.1775	0.2397	0.3236	0.6583	0.8023	0.6526	0.7995	0.0158	0.0164	0.0168	0.0177
LANO (OURS)	0.1437	0.1571	0.1835	0.2405	0.0128	0.0498	0.0934	0.3249	0.0120	0.0125	0.0121	0.0133
Promotion	17.8%	18.2%	20.9%	7.4%	68.4%	26.5%	22.9%	2.2%	45.2%	43.2%	45.7%	51.8%

Table 1: Performance comparison on POBench-PDE under various task settings. Relative L2 error is reported (lower is better). The best results are highlighted in bold, and the second best are underlined. Promotion refers to the relative error reduction w.r.t. the second best model ($1 - \frac{\text{Our error}}{\text{The second best error}}$). “/” denotes that the baseline model is not applicable.

4 Experiment

4.1 POBench-PDE

POBench-PDE is a benchmark suite for PDE-governed tasks under partial observation. It contains six settings across three representative PDE tasks. POBench-PDE is specifically constructed to investigate the central question: *can masked prediction tasks within observed regions implicitly guide operator networks to accurately infer unobserved regions?*

Benchmark Construction. POBench-PDE comprises benchmark PDE solvers and real-world applications, with all datasets reformulated for partially observed trajectories using temporally consistent masks. POBench-PDE includes: • **Navier–Stokes.** 2D turbulent flow (Li et al. 2020). • **Diffusion-Reaction.** Biological pattern formation (Takamoto et al. 2022). • **ERA5.** Real-world climate data (Hersbach et al. 2020).

We exclude tasks like Elasticity that require globally complete inputs, as they do not support meaningful partial observation training. See Appendix C for more details.

Task. To simulate realistic scenarios of partial observation, POBench-PDE supports controlled masking strategies dur-

ing training and evaluation: • **Generalization under Varying Missing Rates.** We simulate varying levels of observational sparsity during training, while keeping the inference stage fixed to full-domain. Specifically, we subsample training data on subsets $\mathcal{D}_o \subset \mathcal{D}$ at different sampling ratios: $s \in \{5\%, 25\%, 50\%\}$. • **Generalization to Missing Patterns.** We evaluate model robustness under two typical missing patterns: (i) point-wise missingness: Independent Bernoulli sampling of observation points. (ii) patch-wise missingness: Structured masking of contiguous spatial blocks to simulate real-world occlusion or sensor failure.

Evaluation. All evaluations are performed under partially observed input conditions. However, we compute relative ℓ_2 error over the entire ground truth output, rather than restricting it to observed regions. This stands in contrast to masked loss settings and highlights the model’s extrapolative generalization across unobserved regions.

4.2 Experimental setting

Baselines. We benchmark against representative neural operators on POBench-PDE: • **MIONet** (Jin, Meng, and Lu




Figure 3: (a) Visualization of boundary-first latent propagation in PhLP. Each row shows different masked trajectory variants of the same historical sequence. From left to right: features evolve from shallow to deep layers, demonstrating progressive reconstruction from partial observations to coherent physical representations. See Appendix D.2 for more visualizations. (b) Effectiveness of Mask-to-Predict Training strategy (MPT) on Diffusion-Reaction dataset. Heatmaps show relative L2 error across different data sizes and missing rates under patch-wise missingness. Left: without MPT, Right: with MPT. Lower is better.

SOURCE	BENCHMARKS	#TIME	#DIM	#MESH
FNO	NAVIER-STOKES	20	2D+TIME	4096
PDE BENCH	DIFFUSION REACTION	20	2D+TIME	4096
ECMWF	ERA5	14	2D+TIME	16200

Table 2: Summary of POBench-PDE.

2022): Multi-input DeepONet with autoregressive architecture for interpolation-free predictions. • **OFormer** (Su et al. 2024): Transformer-based operator with flexible input/output handling. • **CORAL** (Serrano et al. 2023): Neural radiance field-inspired mesh-free PDE solver. • **GNOT** (Hao et al. 2023): Neural operator utilizing linear Transformers. • **IPOT** (Lee and Oh 2024): Attention-based operator with compressed latent space. • **LNO** (Wang and Wang 2024): SOTA neural operator (2024). These models decode solution fields in a single step but often yield blurry and incoherent results due to challenges in capturing spatial dependencies. In contrast, our boundary-first autoregressive framework progressively incorporates spatial structure, improving prediction accuracy.

Implementation. For fair comparison, all models are trained by the AdamW optimizer (Loshchilov 2017) with an initial learning rate of 10^{-3} , scheduled by OneCycleLR (Smith 2018). All baseline models follow their official or widely adopted default configurations. Following standard practice, our model employs 8 layers unless otherwise stated. Notably, all models are trained with the MPT strategy, without which none can converge under partial observation. On the two synthetic datasets, we adopt a default patch-wise missingness with a patch size of 4, while for ERA5, we use a larger patch size of 6. See Appendix D for comprehensive descriptions of implementations.

4.3 Main Results

Benchmark PDE Solvers. As shown in Table 1, LANO achieves near state-of-the-art performance across both bench-

ABLATIONS		#MEM (MB)	#TIME (S/EPOCH)	RELATIVE L2 ↓ NS DR	
#FEATS	1	42.35	261.17	0.5254	0.4613
	8	47.23	279.00	0.2244	0.0071
	16	57.76	281.64	0.1516	0.0073
	32	95.40	282.03	0.1274	0.0092
	64	235.38	424.74	0.1198	0.0603
w/o	BF	54.47	242.94	0.1963	0.4649
	TM	95.18	248.34	0.1459	0.0086
	MPT	95.40	244.50	0.4958	0.4722
TOKEN MIXER	MLP ATTN	95.73	255.16	0.1341	0.0089
		95.40	282.03	0.1274	0.0092

Table 3: Ablations on latent features (#Feats), core components (Boundary-First BF, Token Mixer TM, Mask-to-Predict Training MPT), and token mixer designs. We conduct experiments varying feature numbers, removing components (w/o), and replacing token mixers under a point-wise missingness with 25% missing rate. Memory usage is calculated with batch size 1. NS: Navier-Stokes, DR: Diffusion-Reaction.

marks with all partial observation scenarios. Under patch-wise missingness with less than 50% missing rate, our model consistently outperforms the second-best baseline, with relative error reductions exceeding 17.8%.

The reuse-based variant (LANO, default) significantly outperforms the recalculated one (LANO-S), especially on Diffusion-Reaction with high sparsity and structurally disjoint missing regions. We hypothesize that this is due to the instability of decoder-side attention recalculations, which lacks strong guidance and yields incoherent propagation. In contrast, encoder-generated attention—conditioned on aggregated observations—appears to provide stable priors that better guide decoding.

Real-World Applications. As shown in Table 1, LANO also achieves superior accuracy compared to the baselines. In practical applications such as climate modeling and remote sensing, where missing or partial observations are the norm rather than the exception, the ability of LANO to generalize from incomplete data is crucial.




Figure 4: Performance under varying patch sizes on Navier-Stokes (25% patch-wise missingness). LANO achieves consistent improvements of 18.5%–25.0% over LNO (2024). Remarkably, our experiments demonstrate that even under 25% patch-wise missingness, LANO achieves performance comparable to or exceeding that of FNO (2020), one of the foundational works in neural operator learning.

4.4 Model Analysis

Ablations. We conduct comprehensive ablations on LANO covering feature scaling (#Feats), component removal (w/o), and token mixer replacement. From Table 3: • **Feature scaling:** Optimal at 64 features for NS (77.2% improvement: 0.5254 \rightarrow 0.1198) and 8 for DR, with larger sizes causing overfitting in simpler PDEs. • **Component removal:** All components are essential. Removing the boundary-first framework (BF) or mask-to-predict training (MPT) severely degrades performance (289.2% drop for MPT in NS), while token mixer removal affects complex PDEs more than simple ones. • **Token mixer replacement:** Attention outperforms MLP for NS (5.0% improvement), while MLP suffices for simpler DR, indicating task-dependent optimal complexity.

Effectiveness of Mask-to-Predict Training. We evaluate the effectiveness of MPT on Diffusion-Reaction across varying data sizes (200–800 samples) and missing rates (0%–50%) in a patch-wise missingness setting. Figure 3(b) reveals key observations: • **Significant error reduction:** MPT consistently reduces errors across all configurations, achieving a 84.3% improvement (0.6688 \rightarrow 0.1051) at a 50% missing rate with 200 samples. • **Robustness to data scarcity:** Without MPT, errors exceed 0.48 under high missing rates. MPT maintains stable performance below 0.11 across most configurations. • **Enhanced generalization:** Results confirm that MPT effectively addresses supervision gaps in unobserved regions, enabling robust performance under partial observation.

Visualization of Physics-Aware Latent Propagation. To understand PhLP’s processing of partially observed inputs, we visualize intermediate features across layers for the same trajectory under different missing locations. Figure 3(a) shows progressive evolution from shallow (left) to deep layers (right), revealing key insights: • **Progressive refinement:** Features evolve from sparse fragments to coherent physical structures, confirming boundary-first propagation. • **Physically consistent reconstruction:** Regardless of missing lo-




Figure 5: Evaluation of the model scalability in terms of data size and parameter size. Default setting: 1000 cases, 8 layers.

cations, all input variants converge to similar latent features capturing coherent flow structures, demonstrating PhLP’s robustness to partial observation and ability to infer physically meaningful states.

Performance under Varying Patch Sizes. We evaluate model robustness on Navier-Stokes under patch-wise missingness (25% missing rate) across patch sizes 4–8. Figure 4 shows LANO consistently achieves the lowest relative error across all configurations, with 18.5%–25.0% improvements over the LNO (Wang and Wang 2024) baseline. While most baselines degrade with larger patches, LANO maintains stable performance (1.44–1.66 range), demonstrating effective handling of contiguous missing regions. For fair comparison, all models are configured with 12 layers when evaluated under patch sizes 5–8. FFNO (Tran et al. 2023), a scalable variant of FNO, and LNO-Interp use cubic interpolation for data completion before standard training. MPT-trained models consistently outperform these interpolation-based variants.

Scalability. We analyse LANO scalability by varying training data volume and model depth on the Navier-Stokes dataset under patch-wise missingness with 25% missing rate. The results in Figure 5 show consistent performance gains with larger datasets and models, demonstrating its viability as a foundation model backbone for PDE solving.

5 Conclusion

In this work, we explore neural operator learning from partial observation, addressing a critical challenge in PDE solver deployment. Through systematic experiments on POBench-PDE, we identify two key insights: (i) Mask-to-predict training enables robust learning by creating pseudo-missing regions, and (ii) Boundary-first autoregressive framework is crucial for progressive reconstruction while maintaining physical consistency. We propose LANO with physics-aware latent propagation (PhLP), achieving up to 68% error reduction over existing methods. Our approach demonstrates superior performance on both synthetic benchmarks and real-world climate data. While our work establishes a foundation for practical PDE solving in data-scarce environments, challenges remain for broader deployment. Future work should address these limitations by developing adaptive missing pattern generation strategies and extending the framework to handle irregular geometries and higher-dimensional PDEs.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (2024YFE0202900); the National Natural Science Foundation of China under Grants 62436001, 62406019, 62536001, and 62176020; the Beijing Natural Science Foundation (4244096); the Joint Foundation of the Ministry of Education for Innovation Team (8091B042235); the State Key Laboratory of Rail Traffic Control and Safety (RCS2023K006); and the Talent Fund of Beijing Jiaotong University (2024XKRC075).

References

- Alkin, B.; Fürst, A.; Schmid, S.; Gruber, L.; Holzleitner, M.; and Brandstetter, J. 2024. Universal physics transformers: A framework for efficiently scaling neural operators. *Advances in Neural Information Processing Systems*, 37: 25152–25194.
- Azizzadenesheli, K.; Kovachki, N.; Li, Z.; Liu-Schiaffini, M.; Kossaiji, J.; and Anandkumar, A. 2024. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 1–9.
- Brandstetter, J.; Berg, R. v. d.; Welling, M.; and Gupta, J. K. 2022. Clifford neural layers for PDE modeling. *arXiv preprint arXiv:2209.04934*.
- Brandstetter, J.; Worrall, D. E.; and Welling, M. 2022. Message Passing Neural PDE Solvers. In *International Conference on Learning Representations*.
- Cao, S. 2021. Choose a Transformer: Fourier or Galerkin. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, volume 34.
- Chen, W.; Song, J.; Ren, P.; Subramanian, S.; Morozov, D.; and Mahoney, M. W. 2024. Data-efficient operator learning via unsupervised pretraining and in-context learning. *Advances in Neural Information Processing Systems*, 37: 6213–6245.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Dong, H.; Wang, H.; Liu, H.; Luo, J.; and Wang, J. 2024. Accelerating PDE Data Generation via Differential Operator Action in Solution Space. In *International Conference on Machine Learning*, 11395–11411. PMLR.
- Feng, C.-M.; Yan, Y.; Fu, H.; Chen, L.; and Xu, Y. 2021. Task transformer network for joint MRI reconstruction and super-resolution. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*, 307–317. Springer.
- Gao, W.; and Wang, C. 2023. Active learning based sampling for high-dimensional nonlinear partial differential equations. *Journal of Computational Physics*, 475: 111848.
- Gao, W.; Xu, R.; Deng, Y.; and Liu, Y. 2025. Discretization-invariance? On the Discretization Mismatch Errors in Neural Operators. In *The Thirteenth International Conference on Learning Representations*.
- Gupta, J. K.; and Brandstetter, J. 2022. Towards Multi-spatiotemporal-scale Generalized PDE Modeling. *arXiv preprint arXiv:2209.15616*.
- Hao, Z.; Su, C.; Liu, S.; Berner, J.; Ying, C.; Su, H.; Anandkumar, A.; Song, J.; and Zhu, J. 2024. Dpot: Auto-regressive denoising operator transformer for large-scale pde pre-training. *arXiv preprint arXiv:2403.03542*.
- Hao, Z.; Wang, Z.; Su, H.; Ying, C.; Dong, Y.; Liu, S.; Cheng, Z.; Song, J.; and Zhu, J. 2023. Gnot: A general neural operator transformer for operator learning. In *International Conference on Machine Learning*, 12556–12569. PMLR.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. 2020. The ERA5 global reanalysis. *Quarterly journal of the royal meteorological society*, 146(730): 1999–2049.
- Huang, Z.; Wang, H.; Yang, W.; Tang, M.; Xie, D.; Lin, T.-J.; Zhang, Y.; Xing, W. W.; and He, L. 2025. Self-Attention to Operator Learning-based 3D-IC Thermal Simulation. *arXiv preprint arXiv:2510.15968*.
- Jin, P.; Meng, S.; and Lu, L. 2022. MIONet: Learning multiple-input operators via tensor product. *SIAM Journal on Scientific Computing*, 44(6): A3490–A3514.
- Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; and Yang, L. 2021. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440.
- Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2023. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89): 1–97.
- Lee, S.; and Oh, T. 2024. Inducing Point Operator Transformer: A Flexible and Scalable Architecture for Solving PDEs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 153–161.
- Li, Z.; Huang, D. Z.; Liu, B.; and Anandkumar, A. 2023. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388): 1–26.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2020. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
- Li, Z.; Meidani, K.; and Farimani, A. B. 2022. Transformer for partial differential equations’ operator learning. *arXiv preprint arXiv:2205.13671*.
- Li, Z.; Zheng, H.; Kovachki, N.; Jin, D.; Chen, H.; Liu, B.; Azizzadenesheli, K.; and Anandkumar, A. 2024. Physics-informed neural operator for learning partial differential equations. *ACM/JMS Journal of Data Science*, 1(3): 1–27.
- Liu, G.; Reda, F. A.; Shih, K. J.; Wang, T.-C.; Tao, A.; and Catanzaro, B. 2018. Image inpainting for irregular holes

- using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, 85–100.
- Loshchilov, I. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; and Karniadakis, G. E. 2021. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3): 218–229.
- Luo, J.; Wang, J.; Wang, H.; Geng, Z.; Chen, H.; Kuang, Y.; et al. 2024. Neural Krylov iteration for accelerating linear system solving. *Advances in Neural Information Processing Systems*, 37: 128636–128667.
- Lv, Q.; Liu, T.; and Wang, H. 2025. Exploiting Edited Large Language Models as General Scientific Optimizers. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 5212–5237.
- Morice, C. P.; Kennedy, J. J.; Rayner, N. A.; Winn, J.; Hogan, E.; Killick, R.; Dunn, R.; Osborn, T.; Jones, P.; and Simpson, I. 2021. An updated assessment of near-surface temperature change from 1850: The HadCRUT5 data set. *Journal of Geophysical Research: Atmospheres*, 126(3): e2019JD032361.
- Puzyrev, V. 2019. Deep learning electromagnetic inversion with convolutional neural networks. *Geophysical Journal International*, 218(2): 817–832.
- Serrano, L.; Le Boudec, L.; Kassā Koupaī, A.; Wang, T. X.; Yin, Y.; Vittaut, J.-N.; and Gallinari, P. 2023. Operator learning with neural fields: Tackling pdes on general geometries. *Advances in Neural Information Processing Systems*, 36: 70581–70611.
- Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33: 7462–7473.
- Smith, L. N. 2018. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*.
- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.
- Takamoto, M.; Praditia, T.; Leiteritz, R.; MacKinlay, D.; Alessiani, F.; Pflüger, D.; and Niepert, M. 2022. Pdebench: An extensive benchmark for scientific machine learning. *Advances in Neural Information Processing Systems*, 35: 1596–1611.
- Tran, A.; Mathews, A.; Xie, L.; and Ong, C. S. 2023. Factorized Fourier Neural Operators. In *The Eleventh International Conference on Learning Representations*.
- Van Den Oord, A.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. In *International conference on machine learning*, 1747–1756. PMLR.
- Wang, H.; Hao, Z.; Wang, J.; Geng, Z.; Wang, Z.; Li, B.; and Wu, F. 2024. Accelerating Data Generation for Neural Operators via Krylov Subspace Recycling. In *The Twelfth International Conference on Learning Representations*.
- Wang, H.; Wang, J.; Ma, M.; Shao, H.; and Liu, H. 2025a. SymMaP: Improving Computational Efficiency in Linear Solvers through Symbolic Preconditioning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Wang, H.; Xin, H.; Wang, J.; Yang, X.; Zha, F.; Jiang, Y.; et al. 2025b. Mixture-of-Experts Operator Transformer for Large-Scale PDE Pre-Training. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Wang, H.; Yixuan, J.; Wang, J.; Li, X.; Luo, J.; and huan-shuo dong. 2025c. STNet: Spectral Transformation Network for Solving Operator Eigenvalue Problem. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Wang, T.; and Wang, C. 2024. Latent neural operator for solving forward and inverse pde problems. *arXiv preprint arXiv:2406.03923*.
- Wang, Y.; Ren, S.; Lin, Z.; Han, Y.; Guo, H.; Yang, Z.; Zou, D.; Feng, J.; and Liu, X. 2025d. Parallelized autoregressive visual generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12955–12965.
- Wei, Y.; Huang, Z.; Li, H.; Xing, W. W.; Lin, T.-J.; and He, L. 2025a. Vflow: Discovering optimal agentic workflows for verilog generation. *arXiv preprint arXiv:2504.03723*.
- Wei, Y.; Huang, Z.; Zhao, F.; Feng, Q.; and Xing, W. W. 2025b. MECoT: Markov emotional chain-of-thought for personality-consistent role-playing. In *Findings of the Association for Computational Linguistics: ACL 2025*, 8297–8314.
- Wen, G.; Li, Z.; Azzizadenesheli, K.; Anandkumar, A.; and Benson, S. M. 2022. U-FNO—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163: 104180.
- Wu, H.; Hu, T.; Luo, H.; Wang, J.; and Long, M. 2023. Solving high-dimensional pdes with latent spectral models. *arXiv preprint arXiv:2301.12664*.
- Wu, H.; Luo, H.; Wang, H.; Wang, J.; and Long, M. 2024. Transolver: A fast transformer solver for pdes on general geometries. *arXiv preprint arXiv:2402.02366*.
- Xiao, Z.; Hao, Z.; Lin, B.; Deng, Z.; and Su, H. 2023. Improved operator learning by orthogonal attention. *arXiv preprint arXiv:2310.12487*.
- Zachmanoglou, E. C.; and Thoe, D. W. 1986. *Introduction to partial differential equations with applications*. Courier Corporation.
- Zheng, Y.; Zhang, Q.; Yusifov, A.; and Shi, Y. 2019. Applications of supervised deep learning for seismic interpretation and inversion. *The Leading Edge*, 38(7): 526–533.